

Spiral Toolpaths for High-Speed Machining of 2D Pockets with or without Islands

Mikkel Abrahamsen¹

*Department of Computer Science
University of Copenhagen
Universitetsparken 5
DK-2100 København Ø
Denmark
miab@di.ku.dk*

Abstract

We describe new methods for the construction of spiral toolpaths for high-speed machining. In the simplest case, our method takes a polygon as input and a number $\delta > 0$ and returns a spiral starting at a central point in the polygon, going around towards the boundary while morphing to the shape of the polygon. The spiral consists of linear segments and circular arcs, it is G^1 continuous, it has no self-intersections, and the distance from each point on the spiral to each of the neighboring revolutions is at most δ . Our method has the advantage over previously described methods that it is easily adjustable to the case where there is an island in the polygon to be avoided by the spiral. In that case, the spiral starts at the island and morphs the island to the outer boundary of the polygon. It is shown how to apply that method to make significantly shorter spirals in polygons with no islands than what is obtained by conventional spiral toolpaths. Finally, we show how to make a spiral in a polygon with multiple islands by connecting the islands into one island.

Keywords: Computer-aided manufacturing, CNC machining, High-speed machining, Pocket machining, Spiral toolpath

1. Introduction

A fundamental problem often arising in the CAM industry is to find a suitable toolpath for milling a pocket that is defined by a shape in the plane. A CNC milling machine is programmed to follow the toolpath and thus cutting a cavity with the shape of the given pocket in a solid piece of material. The cutter of the machine can be regarded as a circular disc with radius r , and the task is to find a toolpath in the plane such that the swept volume of the disc, when the disc center is moved along the path, covers the entire pocket. We assume for simplicity that the toolpath is allowed to be anywhere in the pocket.

Some work has been made on spiral toolpaths that morphs a point within the pocket to the boundary of the pocket [2, 3, 10, 11, 12]. The method described by Held and Spielberger [10] yields a toolpath that (i) starts at a user-specified point within the pocket, (ii) ends when the boundary is reached, (iii) makes the cutter remove all material in the pocket, (iv) has no self-intersections, (v) is G^1 continuous², (vi) makes the width of material cut away at most δ at any time, where δ is a user-defined constant called the *stepover*. We must have $\delta < r$, since otherwise some material might not be cut away. See figure 3(b) for an example of such a spiral toolpath. It is the result of an algorithm described in the present paper, but has similar appearance as the spirals described by Held and Spielberger [10].

¹A large part of this work was made while the author worked for Autodesk, Inc.

²A plane curve is G^1 continuous or tangent continuous if there exists a continuous and differentiable parameterization of the curve.

Most traditional toolpath patterns have many places where the cutter does not cut away any new material, for instance in retracts where it is lifted and moved in the air to another place for further machining, or self-intersections of the toolpath, where the tool does not cut away anything new when it visits a place for the second time. That may increase machining time and lead to visible marks on the final product. Spiral toolpaths have the advantage that the cutter is cutting during all of the machining and that, at the same time, the user can control the stepover. Spiral toolpaths are particularly useful when doing high-speed machining, where the rotational speed of the cutter and the speed with which it is moved along the toolpath is higher than in conventional milling. We refer to Held and Spielberger [10] for a more detailed discussion of the benefits of spiral toolpaths compared to various other toolpath patterns and more information on CNC milling in general.

Bieterman and Sandstrom [2] and Huertas-Talón et al. [12] give methods for computing spiral toolpaths by solving elliptic partial differential equation boundary value problems defined on the pocket. However, the methods only work for star-shaped pockets³ [10].

We describe an alternative construction of spirals that also satisfy the previously mentioned properties of the construction of Held and Spielberger [10]. They define a wave that starts at a point in the Voronoi diagram of the pocket at time 0. When the time increases, the wave moves towards the boundary of the pocket in every direction so that at time 1, it reaches the boundary everywhere. The shape of the wave at a certain time represents the area machined at that time. Roughly speaking, the spiral is obtained by traveling around the wave while the time increases from 0 to 1. Our method is similar up to this point, but the way we define the wave is different. The wave defined in [10] is the union of growing disks placed on the Voronoi diagram of the pocket. In our model, the wave is at any time a polygon with its corners on the Voronoi diagram of the pocket. Using that model, we define a spiral consisting of line segments which is at last rounded by circular arcs to get a G^1 continuous curve.

In practice, it is very common that there is one or more islands in the pocket that should be avoided by the cutter, for instance if there are areas of material that should not be machined to the same depth. It is only described by Held and Spielberger [10] how to handle simply-connected pockets, i.e., there must be no islands. In their following paper [11], it is described how one can handle a pocket with an island by connecting it to the boundary with a “bridge”, effectively getting a pocket without islands. Thus, the resulting spiral morphs a point to a shape consisting of the island, the bridge, and the pocket boundary. A big advantage of our method is that it has a natural extension to pockets with one island in the sense that the spiral morphs the shape of the island to the pocket boundary. We exploit the fact that the Voronoi diagram of a pocket with an island consists of exactly one cycle and trees rooted at that cycle to extend our wave model to work in this case. This is our biggest new contribution. We shall demonstrate natural applications of this method to make significantly shorter spirals for pockets with no holes than one could obtain with previously described methods. We also show how to handle pockets with multiple holes.

The paper is based on the author’s experiences while developing the morphed spiral strategy for the AutodeskTM CAM products (HSMWorks, Inventor HSM[®], and Fusion 360[®]). The morphed spiral seems to be quite popular among the users. Due to the abundant number of real-world parts that have been available during the development, we guarantee that it is possible to make an efficient industrial-strength implementation of the algorithms described here.

We use Held’s VRONI library for the computation of Voronoi diagrams [8]. All figures in the paper are automatically generated using our implementation of the algorithms.

The rest of the paper is structured as follows: In section 2, we describe our basic method for making a spiral that morphs a point to the boundary in a simply-connected pocket. Section 3 describes how the method is adapted to a pocket with one island. Using that method, we describe in section 4 an alternative spiral in simply-connected pockets which will be superior to the one from section 2 in many cases. In section 5, we show how to construct a spiral around arbitrarily many islands by first connecting the islands into one island. Finally, we conclude the paper in section 6 by suggesting some future paths of development of spiral toolpaths.

³A polygon is *star-shaped* if there exists a point p in the polygon such that the segment pq is contained in the polygon for every other point q in the polygon.

1.1. Notation and other general conventions

We use zero-based numbering of arrays. For a point $p = (x, y)$, the point $\hat{p} = (-y, x)$ is the counter-clockwise rotation of p . Given two distinct points p and q , pq is the segment between p and q and \overrightarrow{pq} is the half-line that starts at p and contains q . Given a set of points S in the plane, ∂S denotes the boundary of S . In algorithms, we use semicolon to separate different statements written on the same line.

2. Computing a spiral in a pocket without islands

In this section we describe a method to compute a spiral in a given simply-connected 2D pocket \mathcal{P} , see figure 3(b). In practice, the boundary of a pocket is often described by line segments and more advanced pieces of curves, such as like circular arcs, elliptic arcs, and splines. However, it is always possible to use a sufficiently accurate linearization of the input, so we assume for simplicity that \mathcal{P} is a polygon.

Our algorithm first constructs a polyline spiral, see figure 3(a). The polyline spiral must respect the stepover δ , i.e., the distance from every point to the neighboring revolutions and the distance from the outermost revolution to $\partial\mathcal{P}$ is at most δ . In section 2.9 we devise a method for rounding the polyline spiral to get a G^1 continuous spiral consisting of line segments and circular arcs.

The corners of the polyline spiral are points on the edges of the Voronoi diagram of \mathcal{P} , and there is a corner at each intersection point between the spiral and the Voronoi diagram. We only consider the part of the Voronoi diagram inside \mathcal{P} . We have found that we get better results in practice by modifying the Voronoi diagram slightly. We describe these modifications in sections 2.7 and 2.8 to avoid too many technical details here. See figures 1(a)–1(c) for a concrete example of the modifications we make on the Voronoi diagram. Let $\mathcal{VD} = \mathcal{VD}(\mathcal{P})$ be the modified Voronoi diagram of the pocket \mathcal{P} . Like the Voronoi diagram of \mathcal{P} , the modified diagram \mathcal{VD} has the following properties which are necessary and sufficient for the computation of the spiral:

- \mathcal{VD} is a plane tree contained in \mathcal{P} ,
- each leaf of \mathcal{VD} is on the boundary $\partial\mathcal{P}$ of \mathcal{P} ,
- there is at least one leaf of \mathcal{VD} on each corner of \mathcal{P} ,
- all the faces into which \mathcal{VD} divides \mathcal{P} are convex.

2.1. The wave model

We imagine that a wave starts at time $t = 0$ at the point p_0 inside \mathcal{P} . The wave moves out in every direction such that at time $t = 1$, it has exactly the same shape as $\partial\mathcal{P}$. The shape of the wave at a specific time is called a *wavefront*. The wave is growing in the sense that if $0 \leq t_1 \leq t_2 \leq 1$, the wavefront at time t_1 is contained in the wavefront at time t_2 . We choose p_0 as a point in the diagram \mathcal{VD} and consider \mathcal{VD} as a tree rooted at p_0 . We define the time at which the wave hits each node and the speed with which it travels on each edge in \mathcal{VD} . The speed of the wave is always constant or decreasing. Thus, we create a continuous map $\theta : \mathcal{VD} \mapsto [0, 1]$ that assigns a time value between 0 and 1 to each point on \mathcal{VD} . If p is a point moving along a path on \mathcal{VD} from p_0 to any leaf, the value $\theta(p)$ increases monotonically from 0 to 1. For each time $t \in [0, 1]$, the wavefront is a polygon inside \mathcal{P} and the vertices of the wavefront are all the points p on \mathcal{VD} such that $\theta(p) = t$. Note that there is exactly one such point on each path from p_0 to a leaf of \mathcal{VD} for a given $t \in [0, 1]$.

We define a time step $\Delta = 1/r$ for some integer r and compute the wavefront at the times $t \in \{0, \Delta, 2\Delta, \dots, r\Delta\}$, where $r\Delta = 1$, see figure 2. By *wavefront i* , we mean the wavefront at time $i\Delta$. We choose r such that the distance from each point on wavefront i to each of the wavefronts $i - 1$ and $i + 1$ is at most δ when $i > 0$ and $i < r$, respectively. In other words, the *Hausdorff distance* between two neighboring wavefronts is at most δ . Recall that the Hausdorff distance between two sets A and B is $\max\{d(A, B), d(B, A)\}$, where $d(A, B) = \max_{a \in A} \min_{b \in B} \|a - b\|$. For each $i = 1, \dots, r$, we compute a revolution of the polyline spiral by interpolating between the wavefronts $i - 1$ and i . We describe in sections 2.5 and 2.6 how to make the wavefronts and the interpolation such that the stepover is respected between neighboring revolutions.

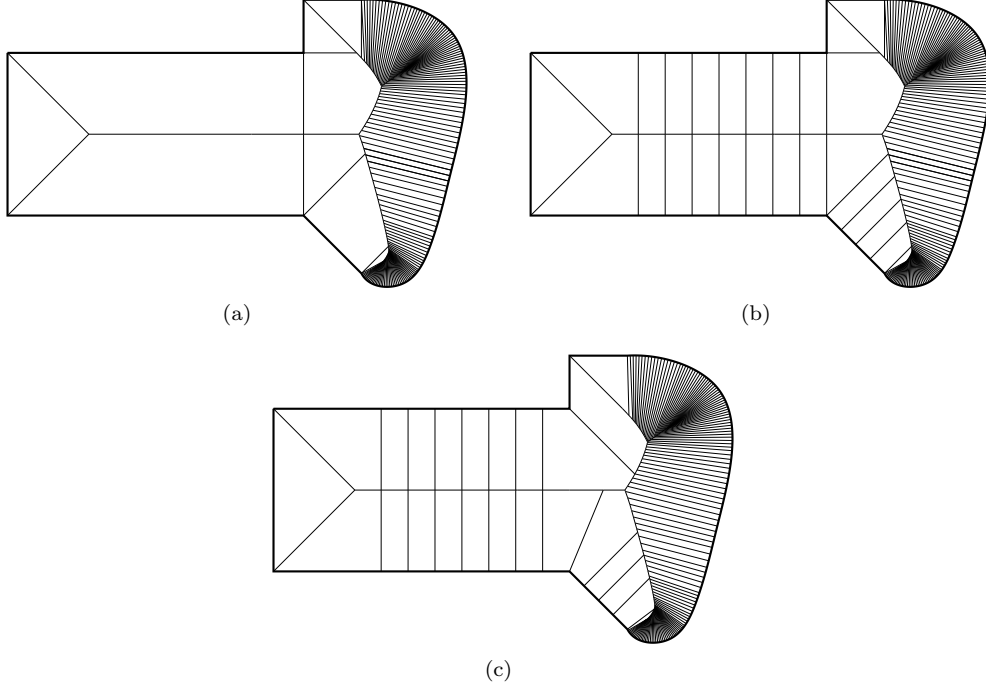


Figure 1: (a) The Voronoi diagram. (b) The Voronoi diagram enriched with equidistantly placed segments perpendicular to long edges. (c) The final diagram $\mathcal{V}\mathcal{D}$ where double edges going to concave corners of \mathcal{P} are replaced by their angle bisector.




Figure 2: The wavefronts in the same polygon \mathcal{P} for the same stepover δ but using two different diagrams to define the wavefronts. (a) The wavefronts obtained using the Voronoi diagram. The Voronoi diagram is in gray. (b) The wavefronts obtained using $\mathcal{V}\mathcal{D}$. $\mathcal{V}\mathcal{D}$ is in gray.




Figure 3: (a) The polyline spiral and \mathcal{VD} in gray. (b) The final rounded spiral.

2.2. Choosing the starting point p_0 and the number of revolutions of the spiral

In order to get a spiral with small length, we try to minimize the number of revolutions. Consider the longest path from p_0 to a leaf in \mathcal{VD} . The length of a path is the sum of edge lengths on the path. If h is the length of the longest path, then $\lceil h/\delta \rceil + 1$ wavefronts are necessary and sufficient for the stepover to be respected between all neighboring wavefronts. Therefore, we choose p_0 as the point in \mathcal{VD} that minimizes the longest distance to a leaf in \mathcal{VD} . That is a unique point traditionally known as the *center* of \mathcal{VD} . Handler [7] gives a simple algorithm to compute p_0 linear time in the size of \mathcal{VD} . The center will most likely not be a node in \mathcal{VD} , but an interior point on some edge. In that case, we split the edge into two edges by introducing a node at p_0 .

2.3. Our representation of \mathcal{VD}

We consider \mathcal{VD} as a directed, rooted tree with the node **Root** at p_0 being the root. We let $\mathbf{Pt}[n]$ be the position of the node n . Let $\mathcal{VD}[n]$ be the subtree rooted at node n . We store a pointer $\mathbf{ParentE}[n]$ to the edge having end node $n \neq \mathbf{Root}$. We say that edge $\mathbf{ParentE}[n]$ is the *parent edge* of node n and any edge having start node n . We also store an array $\mathbf{ChildEs}[n]$ of the edges going out of n sorted in counterclockwise order with the edge following $\mathbf{ParentE}[n]$ being the first. For **Root**, the choice of the first child edge does not matter. For each edge e , we store pointers $\mathbf{Start}[e]$ and $\mathbf{End}[e]$ to the start and end nodes of e . We also store an index $i = \mathbf{IndexInStart}[e]$ such that $\mathbf{ChildEs}[\mathbf{Start}[e]][i] = e$. If e is an edge, we say that $\mathbf{Start}[e]$ and $\mathbf{End}[e]$ are *incident* to e and that e is *incident* to $\mathbf{Start}[e]$ and $\mathbf{End}[e]$. For an edge e_1 and node n incident to e_1 , we let $\mathbf{NextCCW}(e_1, n) = e_2$, where e_2 is the edge after e_1 among the edges incident to n in counterclockwise order. The function $\mathbf{NextCCW}$ can be implemented so that it runs in constant time using the values defined here.

Using $\mathbf{NextCCW}$, we can traverse all of \mathcal{VD} in counterclockwise direction in linear time. We start setting $(n, e) = (\mathbf{Root}, \mathbf{ChildEs}[\mathbf{Root}][0])$. In each iteration, we let n be the other node incident to e and then set $e = \mathbf{NextCCW}(e, n)$. We stop when we have traversed every edge, i.e., when $(e, n) = (\mathbf{Root}, \mathbf{ChildEs}[\mathbf{Root}][0])$ at the end of an iteration. Note that each edge e is visited twice, once going down the tree $\mathcal{VD}[\mathbf{Start}[e]]$ and once going up.

2.4. Defining the movement of the wave

Let $\mathbf{Hgt}[n]$ for each node n be the maximal distance from n to a leaf in $\mathcal{VD}[n]$. All the \mathbf{Hgt} values can be computed in linear time by traversing \mathcal{VD} once. For each node n , we define the time $\mathbf{TmNd}[n]$ where the wave reaches n . We set $\mathbf{TmNd}[\mathbf{Root}] = 0$. We also define the speed $\mathbf{VeNd}[n]$ that the wave has when it reaches n . We set $\mathbf{VeNd}[\mathbf{Root}] = \mathbf{Hgt}[\mathbf{Root}]$. The wave starts at the root at time $t = 0$ and travels with constant speed $\mathbf{VeNd}[\mathbf{Root}]$ on the paths to the farthest leaves in \mathcal{VD} . (Due to our choice of the starting point p_0 , there will always be at least two paths from p_0 to a leaf with maximum length.) Hence, it reaches those leaves at time $t = 1$. On all the shorter paths, we make the wave slow down so that it reaches every leaf at time $t = 1$. In

Appendix A, we describe in detail a possible way of defining the speed and time of the wave on each point in \mathcal{VD} .

When the movement of the wave is defined, we may define $\text{GetPt}(e, t)$ as the point on edge e with $\theta(\text{GetPt}(e, t)) = t$, where $\text{TmNd}[\text{Start}[e]] \leq t \leq \text{TmNd}[\text{End}[e]]$.

2.5. Constructing the wavefronts

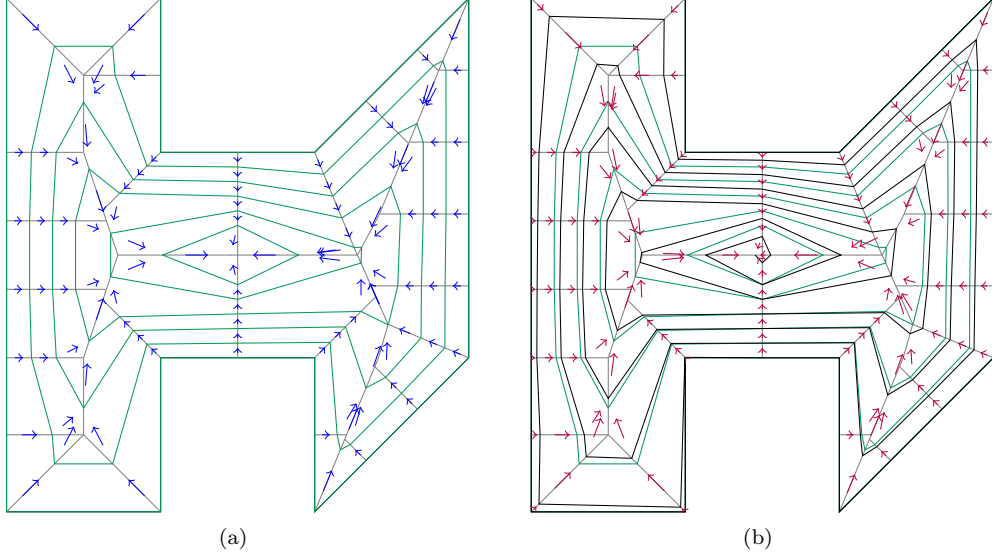


Figure 4: The construction of a polyline spiral in a polygon \mathcal{P} : (a) The wavefronts in green and blue arrows from each wavefront corner $\mathbf{Wf}[i][w]$ to its parent $\mathbf{Wf}[i-1][\text{PaWf}[i][w]]$. The diagram \mathcal{VD} is in gray. (b) The polyline spiral in black obtained by interpolating between the wavefronts. The purple arrows are from each corner $\mathbf{Sp}[i]$ of the spiral to its parent $\mathbf{Sp}[\text{Pa}[i]]$.

We make a spiral with $r = \lceil \frac{\text{Hgt}[\text{Root}]}{\delta'} \rceil$ revolutions, where $\delta' = 0.95 \cdot \delta$. We use the slightly smaller stepover δ' so that the maximum distance between two neighboring revolutions is smaller. That gives more flexibility to smooth the spiral later on as described in section 2.9. We set $\Delta = 1/r$ and compute a wavefront for each of the times $\{0, \Delta, 2\Delta, \dots, r\Delta\}$. The two-dimensional array \mathbf{Wf} stores the wavefronts, so that the wavefront at time $i\Delta$ is the array $\mathbf{Wf}[i]$. Wavefront i is constructed by traversing \mathcal{VD} once and finding every point on \mathcal{VD} with time $i\Delta$ in counterclockwise order. Let e be an edge we have not visited before, and let $n = \text{Start}[e]$ and $m = \text{End}[e]$. There is a corner of wavefront i on e if $\text{TmNd}[n] < i\Delta \leq \text{TmNd}[m]$. If that is the case, we add $\text{GetPt}(e, i\Delta)$ to the end of $\mathbf{Wf}[i]$. We make one corner in $\mathbf{Wf}[0]$ for each of the child edges of the root, $\text{ChildEs}[\text{Root}]$, and all these corners are copies of the point $\text{Pt}[\text{Root}]$. Using this construction, there is exactly one corner of each wavefront on each path from Root to a leaf of \mathcal{VD} .

For each corner $\mathbf{Wf}[i][w]$, we store the length of the part of the wavefront up to the corner, i.e. $\mathbf{WfLng}[i][0] = 0$ and

$$\mathbf{WfLng}[i][w] = \sum_{j=1}^w \|\mathbf{Wf}[i][j] - \mathbf{Wf}[i][j-1]\|,$$

for $w \geq 1$. Here $\|\cdot\|$ is the Euclidean norm. We also store the total length of $\mathbf{Wf}[i]$ as $\text{Tt1WfLng}[i]$.

We introduce a rooted tree with the wavefront corners as the nodes. See figure 4(a). The parent of a corner $\mathbf{Wf}[i][w]$, $i > 0$, is the unique corner $\mathbf{Wf}[i-1][pw]$ on wavefront $i-1$ on the path from $\mathbf{Wf}[i][w]$ to Root . We store pw as $\text{PaWf}[i][w]$, i.e., the parent of $\mathbf{Wf}[i][w]$ is $\mathbf{Wf}[i-1][\text{PaWf}[i][w]]$.

Since the distance from each wavefront corner to each of its children is at most δ' , the Hausdorff distance between two neighboring wavefronts is also at most δ' . Furthermore, since the wave is moving with positive speed towards the leafs of \mathcal{VD} and the faces into which \mathcal{VD} divides \mathcal{P} are convex, neighboring wavefronts do

not intersect each other. From the order in which the corners of a wavefront is constructed, it is also clear that a wavefront does not intersect itself.

2.6. Interpolating between the wavefronts

We construct a polyline spiral stored as an array \mathbf{Sp} . For each $i = 1, \dots, r$, we construct one revolution of the spiral by interpolating between wavefront $i - 1$ and wavefront i . Every corner of the spiral is a point on \mathcal{VD} . There is exactly one spiral corner on the path in \mathcal{VD} from each wavefront corner $\mathbf{Wf}[i][w]$ to its parent $\mathbf{Wf}[i - 1][\mathbf{PaWf}[i][w]]$. Assume for now that we know how to choose the actual corners of the polyline spiral. We shall get back to this shortly.

The first corner $\mathbf{Sp}[0]$ is on the root node of \mathcal{VD} , and for every other corner $\mathbf{Sp}[s]$, $s > 0$, we store an index of the parent $\mathbf{Pa}[s]$, such that the parent $\mathbf{Sp}[\mathbf{Pa}[s]]$ is the first corner we meet on the path in \mathcal{VD} from $\mathbf{Sp}[s]$ to the root. Figure 4(b) shows the resulting polyline spiral and the parents of each corner. We define the spiral such that the distance between a spiral corner and its parent is at most δ' . It follows that the distance from a point on the polyline spiral to the neighboring revolutions is at most δ' .

Here we describe how to define the \mathbf{Pa} -pointers. When we have constructed a spiral corner $\mathbf{Sp}[s]$ which is on the path from $\mathbf{Wf}[i][w]$ to its parent wavefront corner $\mathbf{Wf}[i - 1][pw]$, we know that the first spiral corner on the path from $\mathbf{Wf}[i][w]$ to the root is $\mathbf{Sp}[s]$, and we store this information as $\mathbf{PaSp}[i][w] = s$. Therefore, when we have made a new spiral corner $\mathbf{Sp}[r]$ on the path from $\mathbf{Wf}[i + 1][w']$ to its parent $\mathbf{Wf}[i][pw']$, the parent of $\mathbf{Sp}[r]$ is defined to be $\mathbf{Pa}[r] = \mathbf{PaSp}[i][pw']$. By doing so, the corner $\mathbf{Sp}[\mathbf{Pa}[r]]$ will be the first spiral corner on the path from $\mathbf{Sp}[r]$ to the root.

In the following we describe how to choose the corners of the polyline spiral. We assume that we have finished the revolution of the spiral between wavefronts $i - 2$ and $i - 1$ and we show how to make the revolution between wavefronts $i - 1$ and i . For each wavefront corner $\mathbf{Wf}[i][w]$, we find the point $\mathbf{Q}[w]$ on the path to $\mathbf{Wf}[i - 1][pw]$, where $pw = \mathbf{PaWf}[i][w]$, with time

$$t_w = (i - 1)\Delta + \frac{\mathbf{WfLng}[i][w]}{\mathbf{Tt1WfLng}[i]}\Delta.$$

If $\mathbf{Q}[w]$ is more than δ' away from $\mathbf{Sp}[\mathbf{PaSp}[i - 1][pw]]$, we redefine $\mathbf{Q}[w]$ to be the point on the same path with distance exactly δ' . We mark the path from $\mathbf{Q}[w]$ to the root of \mathcal{VD} . See figure 5.

When we have done the marking for each w , we traverse wavefront i once more. For each wavefront corner $\mathbf{Wf}[i][w]$, we find the first marked point on the path to the root. We let $\mathbf{P}[w]$ be that point and $\mathbf{T}[w] = \theta(\mathbf{P}[w])$ be its time. We have that $\mathbf{T}[w] \geq t_w$, because a later wavefront corner $\mathbf{Wf}[i][w']$, $w' > w$, can mark more of the path from $\mathbf{Wf}[i][w]$ to the root. Therefore, we can have $\mathbf{P}[w] = \mathbf{P}[w + 1]$ for some w . Using this construction, there is exactly one distinct \mathbf{P} -point on each path from a wavefront corner to the root. Furthermore, we know that the distance from $\mathbf{P}[w]$ to the spiral corner $\mathbf{Sp}[\mathbf{PaSp}[i - 1][\mathbf{PaWf}[i][w]]]$ is at most δ' .

The polyline defined by the points $\mathbf{P}[0], \mathbf{P}[1], \dots$ is basically our interpolated spiral, but the points have a tendency to have unnecessarily sharp corners if \mathcal{VD} is relatively dense, which is often the case for polygons \mathcal{P} occurring in real-world problems. To avoid unnecessary sharp corners, we apply a method which we denote as the *convexification*, see figure 5. Let $\mathbf{D}[w] = \sum_{v=0}^{w-1} \|\mathbf{P}[v] - \mathbf{P}[v + 1]\|$ be the length of the polyline until $\mathbf{P}[w]$ and consider the points $(\mathbf{D}[w], \mathbf{T}[w])$. We compute the upper convex hull of these points, e.g. using the method of Graham and Yao [6]. Let \mathcal{F} be the function whose graph is the upper hull. By definition, we have that $\mathbf{T}[w] \leq \mathcal{F}(\mathbf{D}[w])$ for each $w = 0, \dots$. We now choose the corners of \mathbf{Sp} in the following way: For each wavefront corner $\mathbf{Wf}[i][w]$ in order, we find the point S on the path to the root with time $\mathcal{F}(\mathbf{D}[w])$. If S is more than δ' away from the parent spiral corner $\mathbf{Sp}[\mathbf{PaSp}[i - 1][\mathbf{PaWf}[i][w]]]$ (which will be the parent of the spiral corner we are constructing), we choose instead S to be the point on the same path which is exactly δ' away. To avoid repetitions of the same point in \mathbf{Sp} , we add S to the end of \mathbf{Sp} if S is different from the last point in \mathbf{Sp} . Since we get the spiral corners by moving the \mathbf{P} -points closer to wavefront i , we get exactly one distinct spiral corner on each path from a wavefront corner to the root. When \mathcal{VD} is sparse like in figure 4, the convexification makes no visible difference between the \mathbf{P} -points and the final points in \mathbf{Sp} , but when \mathcal{VD} is dense as in figure 5, the effect is significant.




Figure 5: (a) The interpolation between the two blue wavefronts. The boundary \mathcal{P} is in black and the diagram \mathcal{VD} is in gray. The red circles are the points $\mathbf{Q}[w]$ and the marked part of \mathcal{VD} is in black. The green crosses are the resulting points of the polyline spiral stored in \mathbf{Sp} after the convexification process. (b) Related values for the same interpolation: The points $(\mathbf{D}[w], \mathbf{T}[w])$ are red circles. The upper convex hull \mathcal{F} of the points is a black curve, and the green crosses are the points $(\mathbf{D}[w], \mathcal{F}(\mathbf{D}[w]))$ on that hull.

We also add one revolution around \mathcal{P} to the end of \mathbf{Sp} , which is used to test that the last interpolated revolution respects the stepover when the spiral is rounded later on.

The polyline spiral constructed as described clearly satisfies that the distance from a point on one revolution to the neighboring revolutions is at most δ' . Furthermore, each revolution is between two neighboring wavefronts since all the corners of the interpolation between wavefronts i and $i+1$ have times in the interval $[i\Delta, (i+1)\Delta]$. The wavefronts do not intersect as mentioned earlier, so different revolutions of the polyline spiral do not either. It is also clear from the construction that one revolution does not intersect itself. Therefore, the polyline spiral has all the properties that we require of our spiral except being G^1 continuous. How to obtain that is described in section 2.9.

2.7. Enriching the diagram

In this and the following section, we describe some modifications we make on the Voronoi diagram of \mathcal{P} before doing anything else. The result is the diagram $\mathcal{VD} = \mathcal{VD}(\mathcal{P})$. Long edges on \mathcal{P} lead to long faces in the Voronoi diagram, so that the wave is not moving towards the boundary $\partial\mathcal{P}$ in a natural way, see figure 2. In figure 2(a), the wave starts on a long edge, and the first three wavefronts are all degenerated polygons with only two corners, both on the edge. Therefore, we need edges going directly to the boundary with a distance to each other of at most δ , so that each wavefront has corners on more edges than the previous one. We have obtained this by traversing the Voronoi diagram and inspecting each pair of consecutive leafs l_1 and l_2 . Such a pair of nodes are on the same or on two neighboring corners of \mathcal{P} . Assume the latter, so that there is a segment S on $\partial\mathcal{P}$ from l_1 to l_2 and a face f of the Voronoi diagram to the left of S . Let $s = \mathbf{Pt}[l_2] - \mathbf{Pt}[l_1]$ be the vector from l_1 to l_2 , $d = \|s\|$ be the length of S and $m = \lceil d/\delta \rceil$. We want to subdivide f into m faces. Let $p_i = \mathbf{Pt}[l_1] + s \cdot \frac{i}{m}$, $i = 1, \dots, m-1$, be interpolated points on S . Let $h_i = \overrightarrow{p_i, p_i + \hat{s}}$ be the half-line starting at p_i with direction \hat{s} , where \hat{s} is the counterclockwise rotation of s .

For each $i = 1, \dots, m-1$, we find the first intersection point between h_i and the Voronoi diagram. Assume the intersection for some i is a point q on an edge e . If the smallest angle between h_i and e is larger than 50 degrees, we split e into two edges by introducing a node at q and add a segment from that node to a new node at p_i . If the smallest angle is less than 50 degrees, the Voronoi diagram is moving fast enough towards the boundary so that the wavefronts will be fine in that area without adding any additional edges. Figure 1(b) shows the result of enriching the Voronoi diagram shown in figure 1(a).

2.8. Removing double edges to concave corners

A *concave* corner of \mathcal{P} is a corner where the inner angle is more than 180 degrees. Each concave corner c on \mathcal{P} leads to a face in the Voronoi diagram of all the points in \mathcal{P} being closer to c than to anything else on the boundary of \mathcal{P} . Therefore, there are two edges e_1 and e_2 of the Voronoi diagram with an endpoint on c . We have found that we get a better spiral if we remove these edges and instead add an edge following the angle bisector of the edges, i.e., we follow the bisector from c and find the first intersection point q with the Voronoi diagram and add an edge from q to c , see figure 1(c). The reason that this process improves the resulting spiral is that the wavefronts will resemble \mathcal{P} more because they will have one corner on the bisector edge corresponding to the corner c on \mathcal{P} . We can only do this manipulation if the resulting faces are also convex. That is checked easily by computing the new angles of the manipulated faces and it seems to be the case almost always.

2.9. Rounding the polyline spiral

In this section we describe a possible way for smoothing the polyline spiral to get a G^1 continuous spiral. See figure 2.9 for a comparison between the rounded and unrounded spiral from figure 3 using this method. Note that some of the arcs of the rounded spiral rounds multiple corners of the polyline spiral.

For each corner on the polyline spiral, we substitute a part of the spiral containing the corner with a circular arc which is tangential to the polyline spiral in the endpoints. That gives a spiral which is G^1 continuous, i.e., having no sharp corners. Each arc is either clockwise or counterclockwise. For each index i , let $s_i = \text{Sp}[i]\text{Sp}[i+1]$ be the segment from $\text{Sp}[i]$ to $\text{Sp}[i+1]$ and $v_i = \text{Sp}[i+1] - \text{Sp}[i]$ be the vector from $\text{Sp}[i]$ to $\text{Sp}[i+1]$. Each arc has the startpoint p on some segment s_a and the endpoint q on another segment s_b , $a < b$, so that the arc *substitutes* the part of the polyline spiral from p to q . We say that the arc *rounds* the corners $a+1$ to b . We call the arc *tangential* if it is counterclockwise and its center is on the intersection of the half-lines $\overrightarrow{p, p + \vec{v}_s}$ and $\overrightarrow{q, q + \vec{v}_r}$ or it is clockwise and its center is on the intersection of the half-lines $\overrightarrow{p, p - \vec{v}_s}$ and $\overrightarrow{q, q - \vec{v}_r}$.

We store a pointer $\text{Arc}[i]$ to the arc that substitutes the corner $\text{Sp}[i]$. The same arc can substitute multiple consecutive corners, so that $\text{Arc}[i] = \text{Arc}[i+1] = \dots = \text{Arc}[i+k-1]$. In that case, when $\text{Arc}[i-1] \neq \text{Arc}[i] \neq \text{Arc}[i+k]$, the *neighbors* of $\text{Arc}[i]$ are $\text{Arc}[i-1]$ and $\text{Arc}[i+k]$. Two different arcs must substitute disjoint parts of the polyline spiral for the rounded spiral to be well-defined. We subdivide each segment s_a at a point $p_a \in s_a$ such that an arc ending at s_a must have its endpoint at the segment $\text{Sp}[a]p_a$ and an arc beginning at s_a must have its startpoint on the segment $p_a\text{Sp}[a+1]$. The point p_a is chosen as a weighted average of $\text{Sp}[a]$ and $\text{Sp}[a+1]$ so that the arc rounding the sharpest of the corners $\text{Sp}[a]$ and $\text{Sp}[a+1]$ gets most space. Let $\phi_a \in (-\pi, \pi]$ be the angle at corner $\text{Sp}[a]$ of the polyline spiral. We set $w_a = \frac{\pi - \phi_a}{2\pi - \phi_a - \phi_{a+1}}$ and choose p_a as $p_a = (1 - w_a) \cdot \text{Sp}[a] + w_a \cdot \text{Sp}[a+1]$.

We keep a priority queue Q [4] of the arcs that can possibly be enlarged. After each enlargement of an arc, the resulting spiral respects the stepover δ and no self-intersection has been introduced. We say that an arc with these properties is *usable*. Initially, we let each corner be rounded by a degenerated zero-radius arc, and Q contains all these arcs. Clearly, these initial zero-radius arcs are usable by definition. We consider the front arc A in Q and try to find another usable arc A' that substitutes a longer chain of the polyline spiral. The new arc A' normally has a larger radius than A . If possible, we choose A' so that it also substitutes one or, preferably, two of the neighbors of A . Assume that we succeed in making a usable arc A' substituting both A and its neighbors B and C . Now A' rounds the union of the corners previously rounded by A , B , and C . For all these corners, we update the Arc -pointers and we remove A , B , and C from Q . We add A' to Q . For every corner that A' rounds, we also add the arcs rounding the children and parents of that corner

to Q , since it is possible that these arcs can now be enlarged. If no larger usable arc A' is found, we just remove A from Q . The rounding process terminates when Q is empty.

We test if an arc is usable by measuring the distance to the arcs rounding the child and parent corners. That is easily done using elementary geometric computations.

Here we describe how to find the largest usable arc rounding the corners $a + 1$ to b , $a < b$. We find the possible radii of tangential arcs beginning at a point on $p_a \text{Sp}[a + 1]$ and ending at a point on $\text{Sp}[b]p_b$ using elementary geometry. There might be no such arcs, in which case we give up finding an arc rounding these corners. Otherwise, we have an interval $[r_{\min}, r_{\max}]$ of radii of tangential arcs going from $p_a \text{Sp}[a + 1]$ and ending on $\text{Sp}[b]p_b$. We check if the arc with radius r_{\min} is usable. If it is not, we give up. Otherwise, we check if the arc with radius r_{\max} is usable. If it is, we use it. Otherwise, we make a binary search in the interval $[r_{\min}, r_{\max}]$ after the usable arc with the largest radius. We stop when the binary search interval has become sufficiently small, for instance $0.01 \cdot \delta$, and use the arc with the smallest number in the interval as its radius.

The order of the arcs in Q is established in the following way: We have found that giving the arc A the priority $P(A) = r(A)/r(C_{\max}) + 1/s(A)$ gives good results, where $r(A)$ is the radius of A , $s(A)$ is the size of the subtended angle from the center of A in radians, and $r(C_{\max})$ is the radius of the maximum circle contained in \mathcal{P} . The front arc in Q is the one with smallest P -value. We divide by $r(C_{\max})$ to make the rounding invariant when \mathcal{P} and δ are scaled by the same number. $r(C_{\max})$ can be obtained from the Voronoi diagram of \mathcal{P} , since the largest inscribed circle has its center on a node in the diagram. If $s(A) = 0$, we set $P(A) = \infty$, since there is no corner to round.

We see that arcs with small radii or large subtended angles are chosen first for enlargement. In the beginning when all the arcs in Q have zero radius, the arcs in the sharpest corners are chosen first because their degenerated arcs have bigger subtended angles – even though the radius is zero, we can still define the start and end angle of the arc according to the slope of the segments meeting in the corner and thus define the subtended angle of the arc.

If two or three arcs are substituted by one larger arc each time we succeed in making a larger arc, we are sure that the rounding process does terminate, since the complexity of the spiral decreases. However, it is often not possible to merge two or three arcs, but only to make a larger arc rounding the same corners as an old one. The rounded spiral gets better, but we cannot prove that the process terminates. In practice, we have seen fast termination in any tested example. A possible remedy could be only to allow each arc to increase in size without rounding more corners a fixed number of times.

3. Computing a spiral in a pocket with an island

In many practical applications, the area to be machined is not simply-connected, but has one or more “islands” that should not be machined. It might be because there are physical holes in the part or areas of a thicker layer of material not to be machined in the same depth. Therefore, assume that we are given a polygon \mathcal{P} and a single polygonal island \mathcal{I} in the interior of \mathcal{P} . In section 5 we suggest a method to deal with multiple islands. By $\mathcal{P} \setminus \mathcal{I}$, we denote the closed set of points which are in the interior or on the boundary of \mathcal{P} but not in the interior of \mathcal{I} . We want to compute a spiral which is contained in $\mathcal{P} \setminus \mathcal{I}$ such that the Hausdorff distance is at most δ between (i) two consecutive revolutions, (ii) $\partial\mathcal{I}$ and the first revolution, and (iii) $\partial\mathcal{P}$ and the last revolution. As before, δ is the user-defined *stepover*. We also require that the spiral is G^1 continuous and has no self-intersections. See figure 7(b) for an example.

As in the case with a simply-connected pocket, we use a wave model to construct the spiral. We imagine a wave that has exactly the shape of $\partial\mathcal{I}$ at time 0 and moves towards $\partial\mathcal{P}$, so that at the time 1, it reaches $\partial\mathcal{P}$ everywhere. We explain how to define a polyline spiral. The spiral should be rounded by a method similar to the one described in section 2.9.

3.1. The Voronoi diagram of a pocket with an island

We use the Voronoi diagram of the set of line segments of \mathcal{P} and \mathcal{I} . As in the case with no islands, we modify the diagram slightly. Let $\mathcal{VD} = \mathcal{VD}(\mathcal{P} \setminus \mathcal{I})$ be the modified polygon. Like the true Voronoi diagram, the modified diagram \mathcal{VD} has the following properties:




Figure 6: The spirals from figure 3 together. The polyline spiral is blue and the rounded spiral is black.

- \mathcal{VD} is a connected, plane graph contained in $\mathcal{P} \setminus \mathcal{I}$,
- each leaf of \mathcal{VD} is on the boundary of \mathcal{P} or \mathcal{I} ,
- there is at least one leaf of \mathcal{VD} on each corner of \mathcal{P} and \mathcal{I} ,
- all the faces into which \mathcal{VD} divides $\mathcal{P} \setminus \mathcal{I}$ are convex,
- \mathcal{VD} contains exactly one cycle, the cycle is the locus of all points being equally close to $\partial\mathcal{I}$ and $\partial\mathcal{P}$, and \mathcal{I} is contained in its interior.

The diagram \mathcal{VD} is consisting of one cycle \mathcal{C} , which we also denote by the *central cycle*, and some trees growing out from \mathcal{C} , see figure 7(a). Each of the trees grows either outwards and has all its leafs on $\partial\mathcal{P}$ or inwards and has all its leafs on $\partial\mathcal{I}$. The general idea is to use the method described in section 2 to define wavefronts in each of these trees separately. Afterwards, we interpolate between the neighboring wavefronts in each tree and connect the interpolated pieces to get one contiguous spiral.

As in the case of a polygon without an island, we enrich the Voronoi diagram by adding edges equidistantly along and perpendicularly to long edges on $\partial(\mathcal{P} \setminus \mathcal{I})$ as described in section 2.7. One small difference is that if the half-line h_i intersects an edge e at the point q , and e is an edge on the cycle \mathcal{C} , we do not split e and do not introduce a new edge from q to p_i . There would be no advantage of adding that edge. We also remove double edges to concave corners and add their bisector instead as described in section 2.8.

We want the trees to be *symmetric* in the sense that there is a tree \mathcal{PT}_n with root $n \in \mathcal{C}$ and leafs on $\partial\mathcal{P}$ if and only if there is a tree \mathcal{IT}_n with root n and leafs on $\partial\mathcal{I}$. If for a node $n \in \mathcal{C}$ we only have one of the trees, say \mathcal{IT}_n , we add an edge from node n to the closest point on $\partial\mathcal{P}$ and let \mathcal{PT}_n be the tree consisting of that single edge. It follows from the properties of the Voronoi diagram that the added edge does not intersect any of the other edges.

We store \mathcal{C} as a vector $[n_0, \dots, n_{c-1}]$ of the nodes on \mathcal{C} in counterclockwise order, such that there are trees \mathcal{IT}_{n_i} and \mathcal{PT}_{n_i} for each $i = 0, \dots, c-1$. A *root node* is a node n on \mathcal{C} . We let $\mathcal{T}_n = \mathcal{PT}_n \cup \mathcal{IT}_n$ be the union of the two trees rooted at node $n \in \mathcal{C}$ and consider \mathcal{T}_n as a tree rooted at node n .

3.2. Defining the movement of the wave

On each tree \mathcal{T}_n we now define a wave model similar to the one described in section 2.1. The wave starts at time $t = 0$ on the leafs on $\partial\mathcal{I}$ and moves through \mathcal{T}_n so that it hits the leafs on $\partial\mathcal{P}$ at time $t = 1$. Once




Figure 7: (a) A polygon \mathcal{P} with an island \mathcal{I} , both in black. The diagram \mathcal{VD} of $\mathcal{P} \setminus \mathcal{I}$ is drawn with the cycle \mathcal{C} in blue and the other edges in gray. (b) The resulting spiral.

the time and speed in the node n is determined, the times and speeds on the other nodes and edges in \mathcal{T}_n are computed analogously to algorithm 2. One difference is that we compute decreasing times for the tree \mathcal{I}_n . A way to do so is to set $\text{TmNd}[n] = 1 - \text{TmNd}[n]$, compute the inverse times using algorithm 2, and afterwards inverse each of the computed times t by setting $t = 1 - t$.

Let the *preferred time* of a node n be $t_n = \frac{\text{IslHgt}[n]}{\text{BndHgt}[n] + \text{IslHgt}[n]}$, where $\text{BndHgt}[n]$ is the length of the longest path to a leaf in \mathcal{PT}_n and $\text{IslHgt}[n]$ is the length of the longest path to a leaf in \mathcal{IT}_n . Similarly, if the time $\text{TmNd}[n]$ has already been defined, we let the *preferred speed* of n be $v_n = \max \left\{ \frac{\text{IslHgt}[n]}{\text{TmNd}[n]}, \frac{\text{BndHgt}[n]}{1 - \text{TmNd}[n]} \right\}$.

A naive method to define the times and speeds of a root n is to set $\text{TmNd}[n] = t_n$ and $\text{VeNd}[n] = v_n$. That will minimize the number of revolutions and give the most equidistant wavefronts on each tree \mathcal{T}_n . However, the abrupt changes in time and speed along the central cycle \mathcal{C} results in a spiral which curves a lot. Instead, we might smooth the times and speeds around \mathcal{C} . See figure 8 for a comparison of the wavefronts with and without smoothing. There may be many ways of smoothing the values. After much experimentation, we have found the method described in Appendix B to give good results. Notice that the wavefronts in figure 8(b) intersect each other, and that will indeed also be the case for some problem instances after using the smoothing method described in Appendix B. However, in the end of section 3.4 we describe how to ensure that the polyline spiral will not have any self-intersections.

3.3. Creating wavefronts

For a given root node $n \in \mathcal{C}$, we want at least $s_{\mathcal{I}_n} = \frac{\text{IslHgt}[n]}{\delta'}$ revolutions of the spiral in the tree \mathcal{IT}_n in order to respect the stepover $\delta' = 0.95 \cdot \delta$. Similarly, we want $s_{\mathcal{PT}_n} = \frac{\text{BndHgt}[n]}{\delta'}$ revolutions in \mathcal{PT}_n . Therefore, the time between two revolutions should be at most $\Delta_n = \min \left\{ \frac{\text{TmNd}[n]}{s_{\mathcal{I}_n}}, \frac{1 - \text{TmNd}[n]}{s_{\mathcal{PT}_n}} \right\}$. Hence, we let $\Delta' = \min_{n \in \mathcal{C}} \{\Delta_n\}$ be the minimum over all such values. We let the number of revolutions be $r = \lceil 1/\Delta' \rceil$ and set $\Delta = 1/r$.

Each tree \mathcal{T}_n contains a contiguous subset of the corners of any wavefront i . The corners of the subset are points on \mathcal{IT}_n if and only if $t \leq \text{TmNd}[n]$, otherwise they are on \mathcal{PT}_n . Let $r_n = \lfloor \frac{\text{TmNd}[n]}{\Delta} \rfloor$. The wavefronts $i = 0, \dots, r_n$ are on \mathcal{IT}_n , while wavefronts $i = r_n + 1, \dots, r$ are on \mathcal{PT}_n . As explained in section 3.2, we do not use the same time $\text{TmNd}[n]$ for every node n on \mathcal{C} . Therefore, we may get wavefronts crossing \mathcal{C} , as is seen in figure 8.




Figure 8: (a) Wavefronts in red when the times and speeds of the wave in the root nodes have been defined as described in Appendix B. \mathcal{C} is in blue and the rest of \mathcal{VD} is in gray. (b) Wavefronts when the preferred times and speeds have been used to define the movement of the wave.

We suggest to store the wavefront corners in a two-dimensional array for each of the trees \mathcal{IT}_n and \mathcal{PT}_n . The wavefront corners on \mathcal{IT}_n of one wavefront are stored in an array $\text{IslWf}[n][j]$, where index j corresponds to wavefront $i = r_n - j + 1$. In \mathcal{PT}_n , the array $\text{BndWf}[n][j]$ stores corners of wavefront $i = j + r_n$. Hence, the corners of each wavefront is stored locally in each tree \mathcal{T}_n .

In \mathcal{IT}_n , the parents of the corners of wavefront $i = 0, \dots, r_n - 1$ are corners of wavefront $i + 1$. In \mathcal{PT}_n , the parents of the corners of wavefront $i = r_n + 2, \dots, r$ are corners of wavefront $i - 1$. Therefore, all parents are on the wavefront one step closer to the root n . In both \mathcal{IT}_n and \mathcal{PT}_n , we introduce fake wavefront corners at node n stored in the arrays $\text{IslWf}[n][0]$ and $\text{BndWf}[n][0]$, respectively, which are the parents of the corners in the arrays $\text{IslWf}[n][1]$ and $\text{BndWf}[n][1]$. Thus, these fake corners are not corners on wavefront i for any $i = 0, \dots, r$, but are merely made to complete the tree of parent pointers between corners of neighboring wavefronts.

We also need an array WfLng containing *global* information about the length of each wavefront crossing all the trees $\{\mathcal{T}_n\}_n$ in order to do interpolation between the wavefronts later. We have $\text{WfLng}[i][0] = 0$ for every wavefront i . If c_m and c_{m+1} are the m 'th and $(m + 1)$ 'st corners on wavefront i , respectively, we have $\text{WfLng}[i][m + 1] = \text{WfLng}[i][m] + \|c_{m+1} - c_m\|$. Notice that c_m and c_{m+1} can be corners in different, however neighbouring trees \mathcal{T}_n and $\mathcal{T}_{n'}$ and hence stored in different arrays. $\text{TtlWfLng}[i]$ stores the total length of wavefront i .

3.4. Interpolating between wavefronts

We interpolate between two wavefronts $i - 1$ and i in each tree \mathcal{T}_n separately, but using the same technique as in section 2.6. If $i \leq r_n$, we interpolate between the wavefront fragments stored in $\text{IslWf}[n][j]$ and $\text{IslWf}[n][j + 1]$, where $j = r_n - i + 1$, using the values of the length of wavefront $i - 1$ stored in $\text{WfLng}[i - 1]$. If $i > r_n + 1$, we interpolate between $\text{BndWf}[n][j - 1]$ and $\text{BndWf}[n][j]$, where $j = i - r_n$, using the values stored in $\text{WfLng}[i]$. A special case occurs when $i = r_n + 1$, i.e., when we are interpolating between the first wavefront on each side of the root node n . In that case, let

$$t = (i - 1)\Delta + \frac{\text{WfLng}[i][m]}{\text{TtlWfLng}[i]}\Delta,$$

when $\text{BndWf}[n][1][0]$ is the m 'th corner on wavefront i . If $t \leq \text{TmNd}[n]$, we interpolate between $\text{IslWf}[n][0]$ and $\text{IslWf}[n][1]$. Otherwise, we interpolate on the other side of \mathcal{C} , that is, between $\text{BndWf}[n][0]$ and $\text{BndWf}[n][1]$. The convexification process described in section 2.6 can be used in each tree \mathcal{T}_n separately.




Figure 9: A polyline spiral in black and red arrows from each corner $\text{Sp}[s]$ to its parents. The cycle \mathcal{C} is in blue and the diagram \mathcal{VD} is in gray.

We may store the interpolated spiral in an one-dimensional array Sp . Before we add the first interpolated revolution to Sp , i.e., the one between wavefront 0 and 1, we add wavefront 0 to Sp , that is, all the corners on $\partial\mathcal{I}$. Likewise, after the final revolution between wavefronts $r - 1$ and r , we add wavefront r , which is all the corners on $\partial\mathcal{P}$. These are used to ensure that the distance from the first and last revolution to $\partial\mathcal{I}$ and $\partial\mathcal{P}$, respectively, does respect the stepover when rounding the spiral.

For every corner $\text{Sp}[s]$ we have a pointer $\text{Pa}[s]$ such that $\text{Sp}[\text{Pa}[s]]$ is the first spiral corner we meet when traveling from $\text{Sp}[s]$ to the root n of the tree \mathcal{T}_n containing $\text{Sp}[s]$. The parents are not defined for the spiral corners closest to the root node n . Therefore, we make every spiral corner between $\text{IslWf}[n][0]$ and $\text{IslWf}[n][1]$ the parent of every spiral corner between $\text{BndWf}[n][0]$ and $\text{BndWf}[n][1]$ to make parent dependencies across \mathcal{C} . We want all the parent pointers to be towards the island \mathcal{I} . Therefore, we reverse all the pointers between pairs of corners in each tree \mathcal{T}_n . Now, the parent pointers are defined for all spiral corners except for the ones on $\partial\mathcal{I}$. See figure 9 for an example of a polyline spiral around an island and red arrows indicating the parent pointers. Notice that a corner can have multiple parents, but the parents are consecutive and can thus be stored using two indices.

In section 2.6, we stated that since there are no intersections between different wavefronts, the polyline spiral has no self-intersections when there is no island in \mathcal{P} . The wavefronts do not intersect in that case due to the convexity of the faces into which the diagram $\mathcal{VD}(\mathcal{P})$ subdivides \mathcal{P} . When there is an island \mathcal{I} in \mathcal{P} , there are two kinds of faces into which $\mathcal{VD}(\mathcal{P} \setminus \mathcal{I})$ subdivides $\mathcal{P} \setminus \mathcal{I}$, see figure 10. Some faces, like a_3 in the figure, are bounded by edges of one tree \mathcal{T}_n while some are between two trees, like a_1 and a_2 . The latter kind is bounded by edges of two neighboring trees \mathcal{T}_n and \mathcal{T}_m and an edge e from n to m on \mathcal{C} , where n and m are neighboring nodes on \mathcal{C} . The first kind of faces is similar to the faces in section 2, so here we do not worry about self-intersections of the spiral. The second, however, can lead to wavefronts crossing each other and therefore also a self-intersecting spiral as is the case in the figure when the spiral jumps over \mathcal{C} and crosses a_2 . If the union of the faces on each side of e is convex, like a_1 and its neighboring face on the other side of \mathcal{C} , there is no problem. It can easily be tested if the union of the faces is convex by considering the angles of the union at nodes n and m . If it is not, we introduce a new corner on the edge e whenever the spiral jumps from one side of \mathcal{C} to the other. The new corner is an interpolation of nodes n and m using the time of the spiral in the last corner in tree \mathcal{T}_n . Figure 10(c) shows the result of introducing the extra corners. Our experience is that these intersection problems occur very rarely when the method described in Appendix B has been applied.

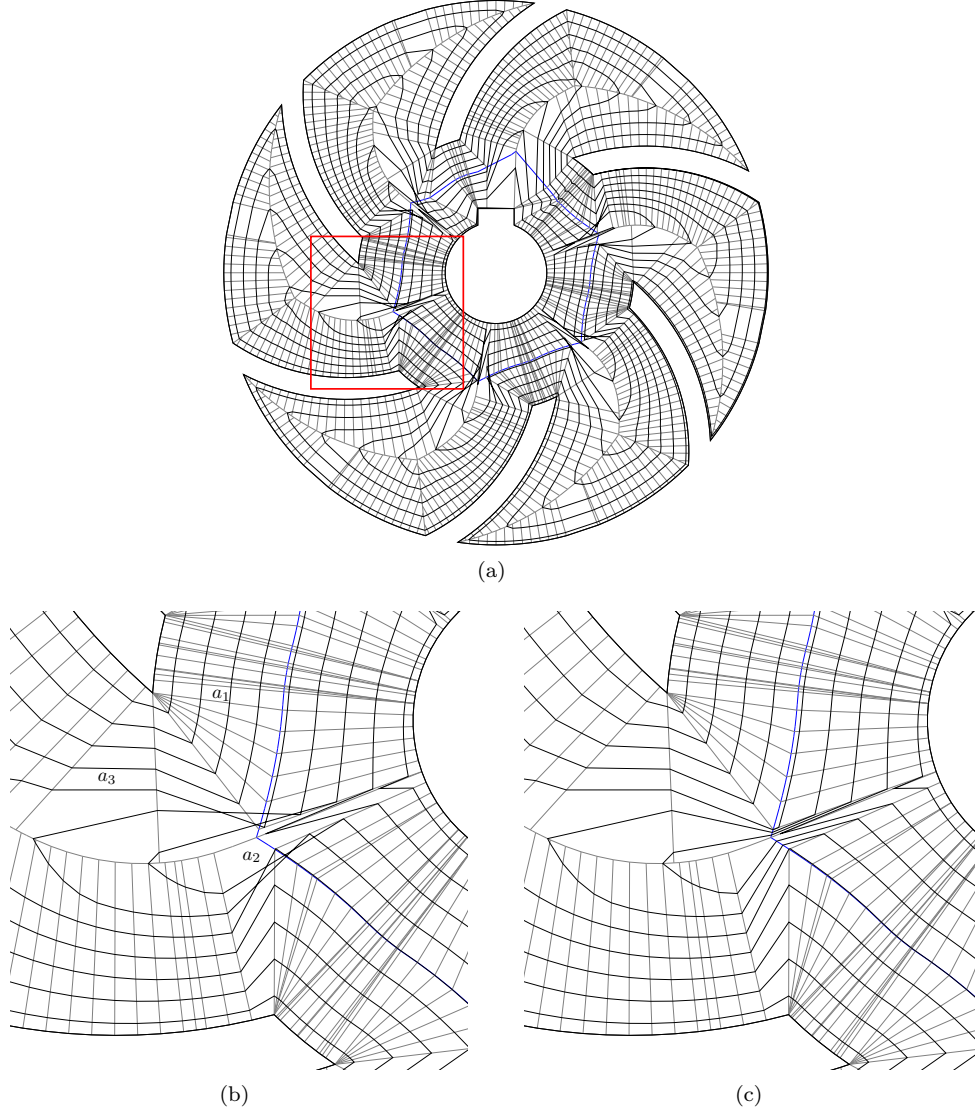


Figure 10: The same pocket and island as in figure 8. The cycle \mathcal{C} is blue, the other edges of \mathcal{VD} are gray. The interpolated spiral is black, and we have not introduced extra spiral corners on \mathcal{C} to avoid self-intersections. In order to emphasize the intersection problems that can arise, we have not used the method described in Appendix B, but merely used the preferred time and speed of each node (section 3.2). (b) is a close-up of (a) of the area in the red rectangle. In (c), we have introduced new corners on \mathcal{C} when the spiral jumps from one side of \mathcal{C} to the other when the union of the two faces on each side of \mathcal{C} is not convex.




Figure 11: Comparison of the basic spiral method from section 2, figure (a), with the improved skeleton method from section 4, figure (b). Note that the spiral obtained from the skeleton method is significantly shorter and that the distance between neighboring revolutions is varying much less than when using the basic method.

4. The skeleton method for a spiral in a pocket without an island

The method from section 2 is mainly applicable if the polygon \mathcal{P} is not too far from being a circle. If \mathcal{P} is very elongated or branched, the distance between neighboring revolutions will often be much less than the maximum stepover. Therefore, the toolpath will be unnecessarily long and the cutting width will vary a lot. That leads to long machining time and an uneven finish of the part. See figure 11 for an example. In such cases, we construct a *skeleton* in \mathcal{P} , which is an island \mathcal{I} with zero area. We then use the method from section 3 to make a spiral from the island to the boundary. It does not matter for the construction of the spiral that the island \mathcal{I} has zero area. See figure 11 for a comparison between the basic method from section 2 with the skeleton method described here when applied to the same polygon.

4.1. Constructing the skeleton in a polygon \mathcal{P}

We choose the skeleton as a connected subset of the edges of the diagram $\mathcal{VD} = \mathcal{VD}(\mathcal{P})$. We traverse \mathcal{VD} once starting at the root and decide for each edge whether to include it in the skeleton. If an edge from node n to m is not included, we don't include anything from the sub-tree $\mathcal{VD}[m]$. For any node n , let $d(n)$ be the length of the shortest path from n to a leaf in $\mathcal{VD}[n]$, and let $D = \max_{n \in \mathcal{VD}} d(n)$. We have found that the following criteria for including an edge e from node n to m gives good results. We require all the criteria to be satisfied.

1. The longest path from n to a leaf in $\mathcal{VD}[n]$ goes through m or $\ell[e] + \text{Hgt}[m] \geq 1.5 \cdot D$, where $\ell[e]$ is the length of edge e .
2. The length of the spanned boundary (defined in Appendix B) of m is larger than $2 \cdot D$.
3. $\text{Hgt}[m] \geq D$.

Criterion 1 is to avoid getting a skeleton that branches into many short paths. Therefore, we only make a branch which is not following the longest path from n if it seems to become at least $0.5 \cdot D$ long (taking criterion 3 into account). When criterion 2 fails, it seems to be a good indicator that an edge is not a significant, central edge in \mathcal{VD} , but merely one going straight to the boundary. Criterion 3 ensures that we do not get too close to the boundary. If we did, we would get very short distances between the neighboring revolutions there. If criterion 3 is the only failing criterion, we find the point p on e such that




Figure 12: (a) The diagram $\mathcal{VD}(\mathcal{P})$ of the polygon \mathcal{P} from figure 11 in green, where the edges chosen for the skeleton are black. (b) The diagram $\mathcal{VD}(\mathcal{P} \setminus \mathcal{I})$ of \mathcal{P} with the skeleton considered an island \mathcal{I} . The cycle \mathcal{C} is blue and the remaining edges are green. The resulting spiral from figure 11(b) is included in gray.

$\text{Hgt}[m] + \|p - \text{Pt}[m]\| = D$ and include the edge from node n to p in the skeleton. See figure 12 for the skeleton constructed given the polygon from figure 11.

If the polygon is close to being a circle, the method described here results in a very small skeleton, and we get a better spiral using the basic method from section 2 in that case. This can be tested automatically by falling back to the basic method if the circumference of the skeleton is less than, say, 5% of the circumference of \mathcal{P} .

5. Computing a spiral in a pocket with multiple islands

The method from section 3 works only for polygons with a single island. If there are many islands $\mathcal{I}_0, \dots, \mathcal{I}_{h-1}$ in a polygon \mathcal{P} , we may connect them with bridges in a tree structure to form one big connected island, see figure 13. The basic idea of reducing the number of islands by connecting them is also used by Chuang and Yang [3] and Held and Spielberg [11]. Our method is given in algorithm 1, and it is a variation of the minimum spanning tree algorithm of Dijkstra [5]. We choose the bridges as edges in the Voronoi diagram V of the area $\mathcal{P} \setminus \bigcup_{i=0}^{h-1} \mathcal{I}_i$. The algorithm creates an array **Bridges** of the edges to use as bridges. We keep a growing set s of the nodes of the Voronoi diagram that we have connected by bridges so far. Here, we have represented s as a bit-vector. We first find one central node n_0 and s is only containing n_0 in the beginning. We use Dijkstra's algorithm [5] in the loop beginning at line 12 to make all shortest paths from nodes in s . We make a relaxation of the distances in line 20. When we reach an island \mathcal{I}_i whose corners are not in s , we use the shortest path to that island as a bridge and add the nodes on the shortest path and the corners on \mathcal{I}_i to s . The use of the distance vector d makes the algorithm prefer to build bridges from the vertices that has been in s for the longest time. That makes the bridges grow from the center node n_0 out in every direction. If we did not choose the bridges in this careful way, the resulting connected island would possibly make unnecessarily long dead ends that would require many revolutions to fill out by the spiral.

VRONI by Held [8] has the possibility to add edges to the Voronoi diagram between neighboring objects in the input where the distance between the objects is shortest, even though these are not genuine Voronoi edges. It is an advantage also to consider these edges when choosing the bridges between the islands, since they are often the best bridges. For instance, all the bridges chosen in the example of figure 13 are of that kind.




Figure 13: (a) A polygon with 13 islands. The bridges chosen among the Voronoi edges to connect the islands are blue. The remaining Voronoi edges are red. (b) The resulting spiral around the islands.

6. Conclusion

We have described methods for the computation of spiral toolpaths suitable for many shapes of pockets for which no previously described algorithms yield equally good results. Our main contribution is the new possibility of making a spiral morphing an island in a pocket to the boundary of the pocket.

Our developed spiral algorithms only work for polygonal input. We believe that it should be possible to generalize to input consisting of line segments and circular arcs, as done by Held and Spielberger [10], using ArcVRONI by Held and Huber [9] to compute the Voronoi diagrams for such input.

Held and Spielberger [11] developed methods to subdivide a pocket with arbitrarily many islands into simply-connected sub-pockets, each of which are suitable for basic spirals as the ones described in section 2. Since we have the possibility to make spirals around islands, we never have to partition the input into separate areas, but in some cases, for instance if the pocket has a long “arm” requiring a lot of revolutions, it might be useful to machine different areas of the input independently. Another possibility is to combine the method for machining around multiple islands with the skeleton method. Given these new possibilities, our problem is quite different from the one discussed by Held and Spielberger. It would be interesting to investigate if and how one can make an automatic partitioning that leads to a more efficient toolpath using our new kinds of spirals.

Acknowledgements

I would like to thank my colleague at Autodesk, Niels Woo-Sang Kjærsgaard, for numerous helpful discussions while I did the development of the spiral algorithms and for many useful comments on this paper. My thanks also go to Autodesk in general for letting me do this research and publish the result.

References

- [1] A. Banerjee, H.-Y. Feng, and E.V. Bordatchev. Process planning for floor machining of $2\frac{1}{2}$ D pockets based on a morphed spiral tool path pattern. *Computers & Industrial Engineering*, 63(4):971–979, 2012.
- [2] M.B. Bieterman and D.R. Sandstrom. A curvilinear tool-path method for pocket machining. *Journal of Manufacturing Science and Engineering, Transactions of the ASME*, 125(4):709–715, 2003.
- [3] J.-J. Chuang and D.C.H. Yang. A laplace-based spiral contouring method for general pocket machining. *The International Journal of Advanced Manufacturing Technology*, 34(7-8):714–723, 2007.

Algorithm 1: MakeBridges($\mathcal{P}, \mathcal{I}_0, \dots, \mathcal{I}_{h-1}$)

```

1  Compute the Voronoi diagram  $V$  of the area inside  $\mathcal{P}$  but outside every  $\mathcal{I}_i$ .
2  Find the set of all nodes  $N$  of  $V$  which are either a corner on an island or equally close to two islands but farther from  $\partial\mathcal{P}$ .
3  Let  $n_0$  be the node in  $N$  closest to the root of  $\mathcal{VD}(\mathcal{P})$  defined as in section 2.2.
4  Let  $s[n] \leftarrow \text{FALSE}$  for each node  $n \in V$ .
5  Let  $d[n] \leftarrow \infty$  for each node  $n \in V$ .
6   $s[n_0] \leftarrow \text{TRUE}$ ;  $d[n_0] \leftarrow 0$ 
7   $\text{Bridges} \leftarrow []$ 
8  while  $s$  does not include every island do
9      For each node  $n$ , if  $s[n] = \text{FALSE}$ , set  $d[n] = \infty$ .
10     Let  $p[n] = \text{NULL}$  for each node  $n$ .
11     Let  $Q$  be a priority queue of the nodes  $n$  of  $V$  sorted after the values  $d[n]$ .
12     while  $Q \neq \emptyset$  do
13         Let  $n$  be the front node in  $Q$  and remove  $n$  from  $Q$ .
14         if  $n$  is a corner on an island  $\mathcal{I}_i$  and  $s[n] = \text{FALSE}$ 
15             For each node  $n'$  on the path  $n, p[n], p[p[n]], \dots$ , set  $s[n'] \leftarrow \text{TRUE}$ .
16             Add each edge on the path  $n, p[n], p[p[n]], \dots$  to  $\text{Bridges}$ .
17             Let  $n_f$  be the node on  $\mathcal{I}_i$  farthest from  $n$  and let  $a \leftarrow \|n_f - n\|$ .
18             For each node  $n'$  on  $\mathcal{I}_i$ , set  $d[n'] \leftarrow d[n] + a$  and  $s[n'] \leftarrow \text{TRUE}$ .
19             Go to line 8.
20     For each edge going out of  $n$  to a node  $n' \in V$ , if  $s[n'] = \text{FALSE}$  and  $d[n] + \|n - n'\| < d[n']$ , set
         $d[n'] \leftarrow d[n] + \|n - n'\|$ , update  $Q$  accordingly, and set  $p[n'] \leftarrow n$ .

```

- [4] T.H. Cormen, C.E. Leiserson, R.L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2009.
- [5] E.W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, 1959.
- [6] R.L. Graham and F.F. Yao. Finding the convex hull of a simple polygon. *Journal of Algorithms*, 4(4):324–331, 1983.
- [7] G.Y. Handler. Minimax location of a facility in an undirected tree graph. *Transportation Science*, 7(3):287–293, 1973.
- [8] M. Held. VRONI: An engineering approach to the reliable and efficient computation of Voronoi diagrams of points and line segments. *Computational Geometry*, 18(2):95–123, 2001.
- [9] M. Held and S. Huber. Topology-oriented incremental computation of Voronoi diagrams of circular arcs and straight-line segments. *Computer-Aided Design*, 41(5):327–338, 2009.
- [10] M. Held and C. Spielberger. A smooth spiral tool path for high speed machining of 2D pockets. *Computer-Aided Design*, 41(7):539–550, 2009.
- [11] M. Held and C. Spielberger. Improved spiral high-speed machining of multiply-connected pockets. *Computer-Aided Design and Applications*, 11(3):346–357, 2014.
- [12] J.L. Huertas-Talón, C. García-Hernández, L. Berges-Muro, and R. Gella-Marín. Obtaining a spiral path for machining STL surfaces using non-deterministic techniques and spherical tool. *Computer-Aided Design*, 50:41–50, 2014.

Appendix A. An algorithm defining the movement of the wave

Here we describe how we define the speed and time of the wave at every point in the diagram $\mathcal{VD}(\mathcal{P})$ in a polygon \mathcal{P} with no island. There might be other models giving equally good or even better results, but we think this model has an advantage of being quite simple to implement. Recall that for a node n in the diagram \mathcal{VD} , $\text{TmNd}[n]$ is the time when the wave reaches n and $\text{VeNd}[n]$ is the speed with which it passes n . We define the movement of the wave from the root and out towards the leafs in \mathcal{VD} . We let $\text{TmNd}[\text{Root}] = 0$ and $\text{VeNd}[\text{Root}] = \text{Hgt}[\text{Root}]$. Let e be an edge going out of the node n where we have defined $\text{TmNd}[n]$ and $\text{VeNd}[n]$ already. Let π be the longest path starting with e . By definition, π has length $h = \ell[e] + \text{Hgt}[\text{End}[e]]$, where $\ell[e]$ is the length of e , i.e., the distance from node n to node $\text{End}[e]$. If $h < \text{Hgt}[n]$, the wave has to slow down on e , since the speed of the wave at node n is determined by $\text{Hgt}[n]$.

We decrease the speed linearly as a function of time such that the wave is decreasing on the first $1/4$ of π while it has constant speed on the last $3/4$ of π . The resulting spiral looks wrong if the wave abruptly changes acceleration when it is not needed. Therefore, if the wave is already slowing down when reaching the node n , we might prefer that it keep slowing down on e with the same rate, even though we must use more than $1/4$ of π . We do that if $\text{Hgt}[n] \leq 1.1 \cdot h$, i.e., if π is almost as long as the longest path going out of n . How to implement this idea is described in greater detail in the following. It would be interesting to find a model where the acceleration is a continuous function of the time along every path in \mathcal{VD} , but we have not found such a model that could be implemented efficiently.

We define the values $\text{TmEd}[e]$ and $\text{VeEd}[e]$ for each edge e which satisfy $\text{TmEd}[e] \geq \text{TmNd}[n]$ and $\text{VeEd}[e] \leq \text{VeNd}[n]$, where n is the node $\text{Start}[e]$. At a time $\text{TmNd}[n] < t < \text{TmEd}[e]$, the speed of the wave is

$$(1 - x) \cdot \text{VeNd}[n] + x \cdot \text{VeEd}[e],$$

where $x = \frac{t - \text{TmNd}[n]}{\text{TmEd}[e] - \text{TmNd}[n]}$. When $t \geq \text{TmEd}[e]$, the speed is $\text{VeEd}[e]$. Let $\text{GetSp}(e, t)$ be the speed defined by the values of edge e and $n = \text{Start}[e]$ at time t . Let $\text{GetDs}(e, t)$ be the distance travelled by the wave from time $\text{TmNd}[n]$ to $t \geq \text{TmNd}[n]$. We have

$$\text{GetDs}(e, t) = \int_{\text{TmNd}[n]}^t \text{GetSp}(e, u) du,$$

which can be computed easily since $\text{GetSp}(e, u)$ is piecewise linear. We also need the function $\text{GetTm}(e, d)$ which is the time t such that $\text{GetDs}(e, t) = d$. Hence, GetTm is the inverse of GetDs , i.e., we have $\text{GetTm}(e, \text{GetDs}(e, t)) = t$ and $\text{GetDs}(e, \text{GetTm}(e, d)) = d$. When $0 < d < \text{GetDs}(\text{TmEd}[e])$, we can compute $\text{GetTm}(e, d)$ by solving a quadratic equation. When $d \geq \text{GetDs}(\text{TmEd}[e])$, we get a linear equation. Finally, $\text{GetPt}(e, t)$ returns the point $(1 - x) \cdot \text{Pt}[n] + x \cdot \text{Pt}[m]$, where $x = \text{GetDs}(e, t) / \ell[e]$ and $m = \text{End}[e]$. For $\text{TmNd}[n] \leq t \leq \text{TmNd}[m]$, $\text{GetPt}(e, t)$ is the position of the wave on edge e at time t .

Assume that we have defined TmNd , VeNd , TmEd , and VeEd on all nodes and edges on the path from Root to some non-leaf node n . Algorithm 2 computes the values for an edge e going out of n and for the node $\text{End}[e]$.

In lines 6–7, we try to use the same values for e as for the previous edge e_p . If, however, the length h of the longest path starting with edge e is smaller than the longest of all paths going out of n , we are in the case of line 9, where we need the wave to slow down. Lines 12–14 compute the distance s that the wave will travel if it continues to decrease speed with the same rate until time 1. We can only keep using the same acceleration if s is smaller than h . If we cannot keep using the same acceleration or the speed of the wave is not decreasing at the node n , we define the values in line 16 as previously described. Both of the lines 16 and 19 give two equations in the two unknowns $\text{TmEd}[e]$ and $\text{VeEd}[e]$. Each pair of equations lead to a quadratic equation in one of the unknowns, and we need to choose the unique meaningful solution. We assign the time and speed values to every node and edge in linear time by traversing \mathcal{VD} once. Algorithm 3 sets times and speeds for all nodes and edges by traversing \mathcal{VD} once.

Algorithm 2: SetTmAndSp(e)

```

1  $n \leftarrow \text{Start}[e]; m \leftarrow \text{End}[e]; t_n \leftarrow \text{TmNd}[n]; v_n \leftarrow \text{VeNd}[n]$ 
2 if  $n = \text{Root}$ 
3    $t_e \leftarrow \text{TmNd}[n]; v_e \leftarrow \text{VeNd}[n]$ 
4 else
5    $e_p \leftarrow \text{ParentE}[n]; t_e \leftarrow \text{TmEd}[e_p]; v_e \leftarrow \text{VeEd}[e_p]$ 
6  $\text{TmEd}[e] \leftarrow t_e$ 
7  $\text{VeEd}[e] \leftarrow v_e$ 
8  $h \leftarrow \text{Hgt}[m] + \ell[e]$ 
9 if  $\text{GetDs}(e, 1) > h$ 
10    $\text{ReuseAcc} \leftarrow \text{False}$ 
11   if  $t_n < t_e$  and  $\text{Hgt}[n] \leq h \cdot 1.1$ 
12      $a \leftarrow \frac{v_e - v_n}{t_e - t_n}$ 
13      $v_1 \leftarrow v_n + a \cdot (1 - t_n)$ 
14      $s \leftarrow \frac{v_n + v_1}{2} \cdot (1 - t_n)$ 
15     if  $s \leq h$ 
16       Define  $\text{TmEd}[e]$  and  $\text{VeEd}[e]$  such that  $\text{GetDs}(e, 1) = h$  and  $\frac{\text{VeEd}[e] - v_n}{\text{TmEd}[e] - t_n} = a$ .
17        $\text{ReuseAcc} \leftarrow \text{True}$ 
18   if  $\text{ReuseAcc} = \text{False}$ 
19     Define  $\text{TmEd}[e]$  and  $\text{VeEd}[e]$  such that  $\text{GetDs}(e, 1) = h$  and  $\text{GetDs}(e, \text{TmEd}[e]) = 0.25 \cdot h$ .
20  $\text{TmNd}[m] \leftarrow \text{GetTm}(e, \ell[e])$ 
21  $\text{VeNd}[m] \leftarrow \text{GetSp}(e, \text{TmNd}[m])$ 

```

Algorithm 3: SetAllTmAndSp()

```

1  $(n_0, e_0) \leftarrow (\text{Root}, \text{ChildEs}[\text{Root}][0])$ 
2  $(n, e) \leftarrow (n_0, e_0)$ 
3 repeat
4   if  $n = \text{Start}[e]$ 
5      $\text{SetTmAndSp}(e)$ 
6      $n \leftarrow \text{End}[e]$ 
7   else
8      $n \leftarrow \text{Start}[e]$ 
9      $e \leftarrow \text{NextCCW}(e, n)$ 
10 until  $(n, e) = (n_0, e_0)$ 

```

Appendix B. A method for defining time and speed on the central cycle

In this section, we describe an alternative method to the naive one mentioned in section 3.2 to define the time and speed of the wave on the cycle \mathcal{C} in the diagram $\mathcal{VD}(\mathcal{P} \setminus \mathcal{I})$ of a polygon \mathcal{P} with an island \mathcal{I} .

Recall that the *preferred time* of a root node n is

$$t_n = \frac{\text{IslHgt}[n]}{\text{BndHgt}[n] + \text{IslHgt}[n]}.$$

We let the time for a node $n \in \mathcal{C}$ be a weighted average of its neighbors' preferred times. We define an *influence distance* $\text{InfDst}[n]$ of each root node $n \in \mathcal{C}$, which is the distance along \mathcal{C} in which n has influence on the times and speeds of other root nodes. In many real-world instances, the majority of the trees \mathcal{T}_n consist of just two edges, namely one going to \mathcal{P} and one to \mathcal{I} , and these two edges are almost equally long. We have experienced that these should have zero influence distance, so that they only have an influence on their own times. Therefore we give a positive influence distance if and only if one of the trees \mathcal{IT}_n and \mathcal{PT}_n have more than one leaf or the ratio $\frac{\text{IslHgt}[n]}{\text{BndHgt}[n]}$ is not in the interval $[1/1.02, 1.02]$.

When the influence distance should be positive, we define it in the following way: Consider three consecutive leafs l_1, l_2 , and l_3 of \mathcal{VD} on $\partial\mathcal{P}$ or $\partial\mathcal{I}$. We define the *spanned boundary* of l_2 to be the path $[M_1, \text{Pt}[l_2], M_2]$, where $M_1 = \frac{\text{Pt}[l_1] + \text{Pt}[l_2]}{2}$ and $M_2 = \frac{\text{Pt}[l_2] + \text{Pt}[l_3]}{2}$. The spanned boundary of a tree \mathcal{PT}_n is the union of all the spanned boundaries of the leafs of \mathcal{PT}_n , similarly for \mathcal{IT}_n . We let $\text{InfDst}[n]$ be the maximum of the distances between the start- and endpoints of the spanned boundaries of \mathcal{PT}_n and \mathcal{IT}_n .

To compute the times of the root nodes, we define a weight of a root node $n \in \mathcal{C}$ as $\mathbf{Wgt}[n] = \mathbf{IslHgt}[n] + \mathbf{BndHgt}[n] + 32 \cdot \mathbf{InfDst}[n]$. If a node n with low weight or zero influence distance is very close (say, closer than 0.1δ) to a node m with a high weight (say, 5 times as much) and positive influence distance, we have experienced better results if we set $\mathbf{Wgt}[n] = 0$. In that way, the influential neighbor m completely dominates node n .

Let n be a fixed node on \mathcal{C} and consider another node m on \mathcal{C} with positive influence distance. Assume that the path from m to n on \mathcal{C} has length $d \leq \mathbf{InfDst}[m]$. We define the weight of node m on node n as $w_m = x_m^3 \cdot \mathbf{Wgt}[m]$, where $x_m = 1 - \frac{d}{\mathbf{InfDst}[m]}$, i.e., we let the weight decrease cubically as the distance increases. The time at node n is defined as

$$\mathbf{TmNd}[n] = \frac{\sum_m w_m t_m}{\sum_m w_m}.$$

Here, the sums are over all nodes m where n is within the influence distance of m .

Recall that the *preferred speed* of a root node n is

$$v_n = \max \left\{ \frac{\mathbf{IslHgt}[n]}{\mathbf{TmNd}[n]}, \frac{\mathbf{BndHgt}[n]}{1 - \mathbf{TmNd}[n]} \right\}.$$

The wave should have a non-increasing speed from a root node n towards the leafs of both \mathcal{IT}_n and \mathcal{PT}_n . Therefore, the speed at node n should at least be the preferred speed so that the wave can reach $\partial\mathcal{I}$ at time 0 and $\partial\mathcal{P}$ at time 1. We define the speed as

$$\mathbf{VNd}[n] = \max_m \left\{ v_m \cdot x_m^2 \cdot (1 - |\mathbf{TmNd}[n] - \mathbf{TmNd}[m]|) \right\},$$

where the maximum is over all the nodes m such that n is within the influence distance of m . The value x_m is defined above. The last factor in the expression is to reduce the influence from nodes that have gotten a very different time than node n , since the speeds of the wave become less comparable when the times are different.